

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE  
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES  
(Attorney Docket № 14222US02)

In the Application of:

Martin Lund

Serial № 10/647,963

Filed: August 26, 2003

For: SYSTEM AND METHOD FOR  
INTEGRATING MULTISERVER  
PLATFORMS

Examiner: BARQADLE, YASIN M

Group Art Unit: 2153

Confirmation № 5243

*Electronically filed on 26-SEP-2008*

**REVISED BRIEF ON APPEAL**

Mail Stop Appeal Brief – Patents  
Commissioner for Patents  
P.O. Box 1450  
Alexandria, VA 22313-1450

Sir:

This Revised Appeal Brief is in response to the Notification of Non-Compliant Appeal Brief mailed on August 29, 2008 (the “Notification”), which indicates that the Summary of Claimed Subject Matter does not give “a concise [sic] explanation of the subject matter defined in each of the independent claims by referring [sic] to the specification by page and line number and to the drawings by reference characters.” The Applicant submits that the previously-filed Appeal Brief’s “Summary of Claimed Subject Matter” section referred to the specification (e.g., “Brief Summary of the

Invention" section) by page and line number and that the referred to page and line numbers concisely explained the subject matter defined in each of the independent claims involved in the appeal as required by 37 CFR 41.37(c)(1)(v). The Applicant notes that it appears that the Examiner is requesting that the independent claims be mapped to the specification, which is **not** required by 37 CFR 41.37(c)(1)(v). Instead, this section states the following:

A concise explanation of the subject matter defined in each of the independent claims involved in the appeal, which shall refer to the specification by page and line number, and to the drawing, if any, by reference characters. For each independent claim involved in the appeal and for each dependent claim argued separately under the provisions of paragraph (c)(1)(vii) of this section, every means plus function and step plus function as permitted by 35 U.S.C. 112, sixth paragraph, must be identified and the structure, material, or acts described in the specification as corresponding to each claimed function must be set forth with reference to the specification by page and line number, and to the drawing, if any, by reference characters.

See 37 CFR 41.37(c)(1)(v). Clearly, the Applicant's previously-filed Appeal Brief's "Summary of Claimed Subject Matter" section concisely explained the subject matter defined in each of the independent claims involved in the appeal. Further, there is nothing in this section that states that the independent claims must be mapped. Nevertheless, the Applicant has mapped the independent claim as requested in the Notification.

The Applicant respectfully requests that the Board of Patent Appeals and Interferences reverse the final rejection of claims 1-15 of the present application. The

Applicants note that this Revised Brief on Appeal is timely because the original Appeal Brief was filed within the two-month period for reply that ended on June 16, 2008 (the Office date of receipt of the Notice of Appeal being April 16, 2008).

**REAL PARTY IN INTEREST**  
**(37 C.F.R. § 41.37(c)(1)(i))**

Broadcom Corporation, a corporation organized under the laws of the state of California, and having a place of business at 5300 California Avenue, Irvine, California 92617, has acquired the entire right, title and interest in and to the invention, the application, and any and all patents to be obtained therefor, as set forth in the Assignment recorded at Reel 014447, Frame 0015 in the PTO Assignment Search room.

**RELATED APPEALS AND INTERFERENCES**  
**(37 C.F.R. § 41.37(c)(1)(ii))**

The Appellant is unaware of any related appeals or interferences.

**STATUS OF THE CLAIMS**  
**(37 C.F.R. § 41.37(c)(1)(iii))**

Claims 1-15 were finally rejected. Pending claims 1-15 are the subject of this appeal.

The present application includes claims 1-15, which are pending in the present application. Claims 1-15 stand rejected under 35 U.S.C. § 102(b) as being anticipated by CISCO Systems, Virtual LAN Communications, <http://web.archive.org/web/19990209172148/cio.cisco.com/warp/public/614/13.html>, July 14, 1995 (hereinafter, CISCO). See Final Office Action at pages 4-11.

The Applicant identifies claims 1-15 as the claims that are being appealed. The text of the pending claims is provided in the Claims Appendix.

**STATUS OF AMENDMENTS**  
**(37 C.F.R. § 41.37(c)(1)(iv))**

The Applicant has not amended any claims subsequent to the final rejection of claims 1-15 mailed on December 14, 2007.

**SUMMARY OF CLAIMED SUBJECT MATTER**  
**(37 C.F.R. § 41.37(c)(1)(v))**

Independent claim 1 recites the following:

1. A method for communication information in a server platform,<sup>1</sup> the method comprising:

receiving at least one packet from at least one of a first switch blade<sup>2</sup> associated with a first multiserver platform<sup>3;4</sup>;

determining at least a server blade<sup>5</sup> associated with a second multiserver platform<sup>6</sup> for receiving at least a portion of said received at least one packet;<sup>7</sup> and

---

<sup>1</sup> See present application, e.g., at page 4, lines 3-4 and page 7, lines 3-4.

<sup>2</sup> See *id.*, e.g., at Figure 1, 160; Figure 2, 202; Figure 3, 306; Figure 4, 408; Figure 6, 607.

<sup>3</sup> See *id.*, e.g., at Figure 1, 100; Figure 2, 201; Figure 3, 303; Figure 4, 402; Figure 6, 604.

<sup>4</sup> See *id.*, e.g., at page 4, lines 4-6 and page 7, lines 4-5.

<sup>5</sup> See *id.*, e.g., at Figure 4, 426.

<sup>6</sup> See *id.*, e.g., at Figure 3, 304; Figure 4, 422; Figure 6, 605.

<sup>7</sup> See *id.*, e.g., at page 4, lines 6-8 and page 7, lines 6-7.

routing said at least a portion of said at least one received packet to at least said server blade<sup>8,9</sup>.

Claims 2-4 are dependent upon claim 1.

Independent claim 5 recites the following:

5. A machine-readable storage having stored thereon, a computer program having at least one code section for communicating information in a server platform, the at least one code section being executable by a machine for causing the machine to perform steps<sup>10</sup> comprising:

receiving at least one packet from at least one of a first switch blade<sup>11</sup> associated with a first multiserver platform<sup>12,13</sup>;

determining at least a server blade<sup>14</sup> associated with a second multiserver platform<sup>15</sup> for receiving at least a portion of said received at least one packet;<sup>16</sup>

and

---

<sup>8</sup> See present application, e.g., at Figure 4, 426.

<sup>9</sup> See *id.*, e.g., at page 4, lines 8-9 and page 7, lines 8-10.

<sup>10</sup> See present application, e.g., at page 4, lines 15-19 and page 19, lines 1-8.

<sup>11</sup> See *id.*, e.g., at Figure 1, 160; Figure 2, 202; Figure 3, 306; Figure 4, 408; Figure 6, 607.

<sup>12</sup> See *id.*, e.g., at Figure 1, 100; Figure 2, 201; Figure 3, 303; Figure 4, 402; Figure 6, 604.

<sup>13</sup> See *id.*, e.g., at page 4, lines 4-6 and page 7, lines 4-5.

<sup>14</sup> See *id.*, e.g., at Figure 4, 426.

<sup>15</sup> See *id.*, e.g., at Figure 3, 304; Figure 4, 422; Figure 6, 605.

<sup>16</sup> See *id.*, e.g., at page 4, lines 6-8 and page 7, lines 6-7.

routing said at least a portion of said at least one received packet to at least said server blade<sup>17,18</sup>.

Claims 6-8 are dependent upon claim 5.

Independent claim 9 recites the following:

9. A system for communicating information in a server platform,<sup>19</sup> the system comprising:

a first multiserver platform<sup>20</sup> comprising at least one of a network interface<sup>21</sup> and a first switch blade<sup>22,23</sup> and  
at least a second multiserver platform<sup>24</sup> comprising a second switch blade<sup>25</sup> coupled<sup>26</sup> to said first switch blade<sup>27</sup> of said first multiserver platform<sup>28,29</sup>.

---

<sup>17</sup> See present application, e.g., at Figure 4, 426.

<sup>18</sup> See *id.*, e.g., at page 4, lines 8-9 and page 7, lines 8-10.

<sup>19</sup> See present application, e.g., at page 4, line 20.

<sup>20</sup> See *id.*, e.g., at Figure 1, 100; Figure 2, 201; Figure 3, 303; Figure 4, 402; Figure 6, 604.

<sup>21</sup> See *id.*, e.g., at Figure 1, 160.

<sup>22</sup> See *id.*, e.g., at Figure 1, 140; Figure 2, 202; Figure 3, 306; Figure 4, 408; Figure 6, 607.

<sup>23</sup> See *id.*, e.g., at page 4, lines 21; page 7, lines 15-19; page 8, lines 24-26; page 9, lines 1-2; page 10, lines 24-25; page 12, lines 2-4; page 15, lines 19-20.

<sup>24</sup> See *id.*, e.g., at Figure 3, 304; Figure 4, 422; Figure 6, 605.

<sup>25</sup> See *id.*, e.g., at Figure 3, 307; Figure 4, 428; Figure 6, 608.

<sup>26</sup> See *id.*, e.g., at Figure 3, 310; Figure 4, 440.

<sup>27</sup> See *id.*, e.g., at Figure 1, 140; Figure 2, 202; Figure 3, 306; Figure 4, 408; Figure 6, 607.

Claims 10-15 are dependent upon claim 9.

**GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL**  
**(37 C.F.R. § 41.37(c)(1)(vi))**

Claims 1-15 stand rejected under 35 U.S.C. § 102(b) as being anticipated by CISCO Systems, Virtual LAN Communications, <http://web.archive.org/web/19990209172148/cio.cisco.com/warp/public/614/13.html>, July 14, 1995 (hereinafter, CISCO).

---

<sup>28</sup> See present application, e.g., at Figure 1, 100; Figure 2, 201; Figure 3, 303; Figure 4, 402; Figure 6, 604.

<sup>29</sup> See *id.*, e.g., at page 4, lines 22-23; page 11, lines 3-5; page 12, lines 13-14.

**ARGUMENT**  
**(37 C.F.R. § 41.37(c)(1)(vii))**

In the Final Office Action, claims 1-15 stand rejected under 35 U.S.C. § 102(b) as being anticipated by CISCO.

**I. Claims 1-15 Are Not Anticipated by CISCO**

Claims 1-15 stand rejected under 35 U.S.C. § 102(b) as being anticipated by CISCO.

**A. Rejection of Independent Claims 1 and 5**

The Applicant turns to the rejection of claims 1 and 5 under 35 U.S.C. § 102(b) as being anticipated by CISCO. The Applicant submits that CISCO does not disclose or suggest at least the limitations of “receiving at least one packet from at least one of a **first switch blade associated with a first multiserver platform**; determining at least a **server blade associated with a second multiserver platform** for receiving at least a portion of said received at least one packet; and routing said at least a portion of said at least one received packet to at least **said server blade**,” as recited by the Applicant in the independent claim 1 (emphasis added).

With regard to “receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform,” the Final Office Action states “see Page 9, Figure 9, left hand side bottom switch is first blade switch and multiserver platform is one under label VLAN1; it is inherent with network interface that any packet

originated by the multiserver will be received by the first blade switch." See Final Office Action, Page 4, Lines 12-16. **However, CISCO fails to disclose a first switch blade or a first multiserver platform. Rather, the CISCO publication discloses using the Catalyst 5000, Catalyst 1200 and ProStack switches, which are not switch blades. By explicitly teaching to use non-blade switches, CISCO clearly teaches away from using "a first switch blade" as set forth in Applicant's independent claims 1 and 5.**

Additionally, the Applicant notes that claims 1 and 5 recite "receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform." Thus, the Examiner's assertion that "it is inherent with network interface that any packet originated by the multiserver will be received by the first blade switch" is irrelevant. The Final Office Action fails to show where CISCO discloses "receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform," as set forth in Applicant's independent claims 1 and 5.

The Applicant further notes that the only servers disclosed in CISCO's Fig. 9 are in the bottom right hand corner. The computer icon is used in the CISCO publication to show "end stations" and the file cabinet is used to show a "server." See e.g., Figs. 2, 5, 7 and 9. The Examiner acknowledged the difference between workstations (i.e., the computer icon) and servers (i.e., the file cabinet icon) in the Response to Arguments section of the Final Office Action where the Examiner stated that "VLANs as in fig. 2 show a multiple VLAN groups of an enterprise network that include Engineering VLAN, Marketing VLAN and accounting VLAN each with its own workstations, servers

connected via switch devices." See Final Office Action, Page 3, Lines 10-13 (emphasis added). Clearly, at least Fig. 9 of the CISCO publication differentiates between "end stations" and "servers." Thus, the computer icon under VLAN1 in the bottom left hand corner of Fig. 9 cannot be a multiserver platform as alleged by the Examiner. Further, nothing in CISCO indicates that the VLAN1, VLAN2, or VLAN3 servers in the bottom right hand corner of CISCO's Fig. 9 are multiserver platforms or part of a multiserver platform. A multiserver platform is described in Applicant's specification in, for example, at least Figure 1 and accompanying text in paragraphs 22-28.

With regard to "determining at least a server blade associated with a second multiserver platform for receiving at least a portion of said received at least one packet," the Final Office Action states the following:

[S]ee Page 9, Figure 9, The second blade switch is left top switch and multiserver platform is the one under VLAN1; when a packet is received from first multiserver platform is received by the switch, the switch will determine (determination is made by rules set by administrator) if the packet is to be sent to second multiserver platform, "Both of these techniques examine the packet when it is either received or forwarded by the switch. Based on set of rules defined by the administrator...", Page 3, third Para, lines 4-6).

See Final Office Action, Page 4, Line 19 – Page 5, Line 6. However, nowhere in CISCO is there any disclosure regarding a server blade and the Final Office Action fails to mention a server blade. It appears that the Final Office Action confuses a "server blade" with a "second blade switch." A server blade is different than a switch blade.

(See e.g., Applicant's Figure 1 Blade Server 120 compared to Switch Blade 140 and accompanying text in paragraphs 25-27). Further, as discussed above, nowhere in CISCO is there any disclosure regarding a switch blade or a multiserver platform.

With regard to "routing at least a portion of said at least one received packet to at least said server blade," the Office Action states the following:

**Once the determination is made by the first blade switch, that packet belongs to multiserver platform connected to second switch blade, it will be sent (routed to second blade that is associated with second multiserver platform, "Based on the set of rules defined by the administrator, these techniques determine where the packet is to be sent, filtered, and/or broadcast.", Page 3, third Para, Lines 5-7. See also page 9, first paragraph for end-user VLAN information and identification carried between switches, routers, and directly attached servers.**

See Final Office Action, Page 5, Lines 8-17 (emphasis added). The Examiner asserts that the packet is routed from a first blade switch to a second switch blade associated with a second multiserver platform. However, as mentioned above, it appears that the Final Office Action confuses a "server blade" with a "second switch blade." A server blade is different than a switch blade. See e.g., Figure 1 Blade Server 120 compared to Switch Blade 140 and accompanying text in paragraphs 25-27. Nowhere in CISCO is there any mention of a server blade, let alone "routing said at least a portion of said at least one received packet to at least said server blade," as set forth in Applicant's independent claims 1 and 5.

Accordingly, independent claims 1 and 5 are not anticipated by CISCO and are allowable. Furthermore, the Applicant reserves the right to argue additional reasons beyond those set forth herein to support the allowability of claims 1 and 5.

## B. Examiner's Response to Arguments

The Examiner states the following in the Final Office Action:

Examiner notes that switch blade is a device that performs switching functions (normally at layer 2 of the OSI model). ... As such the switches in Cisco document meet the functions performed by the Applicant's claimed switch blades.

See Final Office Action, Page 2, Lines 10-11 and 13-15. The Applicant points out that with regard to the anticipation rejections, MPEP 2131 states, "[a] claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference." *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631 2 USPQ2d 1051, 1053 (Fed.Cir. 1987) (emphasis added). MPEP 2131 also states, "[t]he identical invention must be shown in as complete detail as is contained in the ... claim." *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989) (emphasis added).

The Examiner does not contest that CISCO is silent as to a switch blade. Rather, the Examiner asserts that "the switches in Cisco document meet the functions performed by the Applicant's claimed switch blades." Thus, using the Examiner's

reasoning with regard to, for example, “receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform,” any device that sends at least one packet could be considered “a first switch blade” because that device would allegedly “meet the functions performed by the Applicant’s claimed switch blades.” Clearly, CISCO fails to disclose each and every element as set forth in Applicant’s independent claims 1 and 5, as required by MPEP 2131 because CISCO at least fails to disclose “a first switch blade.” Further, CISCO fails to teach the identical invention in as complete detail as is contained in independent claims 1 and 5, as required by MPEP 2131 because CISCO at least fails to disclose “a first switch blade.”

Further, the Applicant notes that “switch blade” is a known term in the art. As already stated above, CISCO does not disclose the use of a “switch blade.” Even a broad interpretation of CISCO cannot overcome at least this deficiency.

The Examiner further states the following in the Final Office Action:

As to the issue of server blades and the multiserver platform, the Cisco reference teaches “The backbone commonly acts as the aggregation point for large volumes of traffic. It also carries end-user VLAN information and identification between switches, routers, and directly attached servers.” (Page 9, first paragraph. See also *Figure 7: Servers as Part of Multiple VLANs*). It is also noted that VLANs as in fig. 2 show a multiple VLAN groups of an enterprise network that include Engineering VLAN, Marketing VLAN and accounting VLAN each with its own workstations, servers connected via switch devices.

See Final Office Action, Page 3, Lines 4-13 (emphasis in original). With regard to “a server blade,” the Applicant notes that nowhere in the Examiner’s Response to Arguments section of the Final Office Action is there any identification of “a server

blade" in CISCO. Rather, the Response to Arguments section of the Final Office Action merely points out that CISCO discloses servers. Nowhere in CISCO is there any mention of "a server blade." Thus, CISCO clearly fails to disclose each and every element as set forth in Applicant's independent claims 1 and 5, as required by MPEP 2131 because CISCO at least fails to disclose "a server blade." Further, CISCO fails to teach the identical invention in as complete detail as is contained in independent claims 1 and 5, as required by MPEP 2131 because CISCO at least fails to disclose "a server blade."

Further, the Applicant notes that "server blade" is a known term in the art. As already stated above, CISCO does not disclose the use of a "server blade." Even a broad interpretation of CISCO cannot overcome at least this deficiency.

With regard to "a multiserver platform," the Applicant notes that the Examiner's Response to Arguments section of the Final Office Action merely discusses CISCO's disclosure of using more than one server. However, the use of more than one server does not necessarily mean using a multiserver platform. For example, the Applicant's specification discusses the disadvantages of using multiple single servers in at least paragraphs [06] and [07]. Nowhere in CISCO is there any mention of using multiserver platforms. Thus, CISCO clearly fails to disclose each and every element as set forth in Applicant's independent claims 1 and 5, as required by MPEP 2131 because CISCO at least fails to disclose "a multiserver platform." Further, CISCO fails to teach the identical invention in as complete detail as is contained in independent claims 1 and 5,

as required by MPEP 2131 because CISCO at least fails to disclose “a multiserver platform.”

Further, the Applicant notes that “multiserver platform” is a known term in the art. As already stated above, CISCO does not disclose the use of a “multiserver platform.” Even a broad interpretation of CISCO cannot overcome at least this deficiency.

### **C. Rejection of Dependent Claims 2 and 6**

Claims 2 and 6 depend on independent claims 1 and 5, respectively. Therefore, the Applicant submits that claims 2 and 6 are allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 1 and 5. The Applicant also submits that CISCO does not disclose or suggest at least the limitation of “wherein said receiving further comprises receiving said at least one packet by at least one of a second switch blade associated with a third multiserver platform and a central switch,” as recited by the Applicant in claims 2 and 6.

With regard to claims 2 and 6, the Final Office Action states the following at page 3:

The method and computer according to claim 1 and 5 (see supra for discussion of claims 1 and 5), wherein said receiving further comprises receiving said at least one packet by at least one of a second switch blade and central switch (see figs 3; 5 and Figure 9) it is inherent in the FDDI a packet on the ring will be received by all members of the ring.

See Final Office Action, Page 5, Line 19 – Page 6, Line 2. For reasons similar to those discussed above with regard to claims 1 and 5, CISCO fails to disclose “a second

switch blade" and "a third multiserver platform." Additionally, nowhere in CISCO is there any mention of "a central switch." While discussing claims 3 and 7, the Examiner states in the Final Office Action that "[s]ee page 9, Figure 9, the central switch is middle top switch on FDDI ring and is connected through FDDI ring to the bottom left (first blade switch) and top left (second blade switch) switch." See Final Office Action, Page 6, Lines 11-14. However, the icon referred to by the Examiner in Figure 9 is a router, not a central switch as alleged by the Examiner. For example, CISCO's Figure 2 shows the same icon and labels it "Cisco Router." Additionally, CISCO's Figure 7 shows the same icon and labels it "Cisco 7000." The Applicant submits that a router is different than a central switch. Thus, CISCO fails to disclose "a central switch" as recited in Applicant's dependent claims 2 and 6.

Accordingly, the Applicant submits that claims 2 and 6 are allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claims 2 and 6.

#### **D. Rejection of Dependent Claims 3 and 7**

Claims 3 and 7 depend on dependent claims 2 and 6, respectively; which depend from independent claims 1 and 5, respectively. Therefore, the Applicant submits that claims 3 and 7 are allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 1, 2, 5 and 6.

Accordingly, the Applicant submits that claims 3 and 7 are allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claims 3 and 7.

**E. Rejection of Dependent Claims 4 and 8**

Claims 4 and 8 depend on independent claims 1 and 5, respectively. Therefore, the Applicant submits that claims 4 and 8 are allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 1 and 5.

Accordingly, the Applicant submits that claims 4 and 8 are allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claims 4 and 8.

**F. Rejection of Independent Claim 9**

The Applicant turns to the rejection of claim 9 under 35 U.S.C. § 102(b) as being anticipated by CISCO. The Applicant submits that CISCO does not disclose or suggest at least the limitation of **“a first multiserver platform comprising at least one of a network interface and a first switch blade; and at least a second multiserver platform comprising a second switch blade coupled to said first switch blade of said**

first multiserver platform," as recited by the Applicant in the independent claim 9 (emphasis added).

With regard to "a first multiserver platform comprising at least one of a network interface and a first switch blade," the Final Office Action states the following:

[A] first multiserver platform (for this claim, Page 8, Figure 8, right hand side block will be used for claim elements; bottom row) comprising at least one of a network interface (multiserver platform first one in bottom row, under VLAN1 and a network interface connecting to the switch to the multiserver) and a first switch blade (bottom block of three switches, second one connected to first multiserver platform under the label VLAN1).

See Final Office Action, Page 7, Lines 9-16. As mentioned above with regard to claims 1 and 5, nowhere in CISCO is a multiserver platform disclosed. CISCO's VLAN1 in Figure 8 is different than a multiserver platform. A multiserver platform is disclosed in Applicant's specification in, for example, at least Figure 1 and accompanying text in paragraphs 22-28. Further, neither Figure 8 nor the supporting text in CISCO mentions a network interface. CISCO discloses workstations connected to switches; however, because none of the workstations are shown directly connected to the network, CISCO can not disclose a workstation comprising a network interface, let alone a multiserver platform comprising a network interface. See CISCO, Figure 8, right hand side. Additionally, as mentioned above with regard to claims 1 and 5, nowhere in the CISCO reference is there any mention of using switch blades or server blades. Rather, the CISCO reference discloses using the Catalyst 5000 and ProStack switches, which are not switch blades.

Also, Applicant's claim 9 recites "a first multiserver platform comprising at least one of a network interface and a first switch blade." The Examiner asserts that the VLAN1 workstation in the bottom row of the right hand block in Figure 8 is a first multiserver platform (which it is not) and that it is connected to the ProStack switch in the bottom row of the right hand block in Figure 8. Thus, even if the workstation was a first multiserver platform (which it is not), CISCO fails to show the workstation comprising at least one of a network interface and a first switch blade. Rather, CISCO's Figure 8 shows the workstation connected to the switch. Nowhere in CISCO is there any disclosure that the VLAN workstations comprise at least one of a network interface and a first switch blade.

With regard to "at least a second multiserver platform comprising a second switch blade coupled to said first switch blade of said first multiserver platform," the Final Office Action states the following:

[A]t least a second multiserver platform (second row) comprising a second switch blade (middle switch) coupled [to] said first switch blade of said first multiserver platform (middle switch connected to second multiserver platform, under the label VLAN1; both first multiserver platform and second multiserver platform are coupled by VLAN1).

See Final Office Action, Page 7, Line 16 - Page 8, Line 2. However, as mentioned above, nowhere in CISCO is a multiserver platform disclosed. CISCO's VLAN1 in Figure 8 is different than a multiserver platform. A multiserver platform is disclosed in Applicant's specification in, for example, at least Figure 1 and accompanying text in paragraphs 22-28. Further, as mentioned above, nowhere in the CISCO reference is

there any mention of using switch blades. Rather, the CISCO reference discloses using the Catalyst 5000 and ProStack switches, which are not switch blades.

Additionally, Applicant's claim 9 recites "at least a second multiserver platform comprising a second switch blade coupled to said first switch blade of said first multiserver platform." The Examiner asserts that CISCO discloses "at least a second multiserver platform (second row) comprising a second switch blade (middle switch)...." See Final Office Action, Page 7, Lines 16-18. However, the CISCO's VLAN1 workstation in the second row of the right hand block in Figure 8 is shown as being connected to the Catalyst 5000 switch in the middle row of the right hand block in Figure 8. Thus, even if the workstation was a second multiserver platform (which it is not), CISCO fails to show the workstation comprising a second switch blade. Rather, CISCO's Figure 8 shows the workstation connected to the switch. Nowhere in CISCO is there any disclosure that the VLAN workstations comprise a second switch blade.

Accordingly, independent claim 9 is not anticipated by CISCO and is allowable. Furthermore, the Applicant reserves the right to argue additional reasons beyond those set forth herein to support the allowability of claim 9.

## **G. Examiner's Response to Arguments**

The Examiner states the following in the Final Office Action:

Examiner notes that switch blade is a device that performs switching functions (normally at layer 2 of the OSI model). ... As such the switches

in Cisco document meet the functions performed by the Applicant's claimed switch blades.

See Final Office Action, Page 2, Lines 10-11 and 13-15. The Applicant points out that with regard to the anticipation rejections, MPEP 2131 states, “[a] claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference.” *Verdegaal Bros. v. Union Oil Co. of California*, 814 F.2d 628, 631 2 USPQ2d 1051, 1053 (Fed.Cir. 1987) (emphasis added). MPEP 2131 also states, “[t]he identical invention must be shown in as complete detail as is contained in the ... claim.” *Richardson v. Suzuki Motor Co.*, 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989) (emphasis added).

The Examiner does not contest that CISCO is silent as to a switch blade. Rather, the Examiner asserts that “the switches in Cisco document meet the functions performed by the Applicant's claimed switch blades.” However, Applicant's independent claim 9 is a system claim; and therefore, components of Applicant's system are claimed and CISCO fails to disclose the claimed components of the system. Clearly, CISCO fails to disclose each and every element as set forth in Applicant's independent claim 9, as required by MPEP 2131 because CISCO at least fails to disclose “a first switch blade” and “a second switch blade.” Further, CISCO fails to teach the identical invention in as complete detail as is contained in independent claim 9, as required by MPEP 2131 because CISCO at least fails to disclose “a first switch blade” and “a second switch blade.”

Further, the Applicant notes that “switch blade” is a known term in the art. As already stated above, CISCO does not disclose the use of a “switch blade.” Even a broad interpretation of CISCO cannot overcome at least this deficiency.

With regard to “a multiserver platform,” the Applicant notes that the Examiner’s Response to Arguments section of the Final Office Action merely discusses CISCO’s disclosure of using more than one server. However, the use of more than one server does not necessarily mean using a multiserver platform. For example, the Applicant’s specification discusses the disadvantages of using multiple single servers in at least paragraphs [06] and [07]. Nowhere in CISCO is there any mention of using multiserver platforms. Thus, CISCO clearly fails to disclose each and every element as set forth in Applicant’s independent claim 9, as required by MPEP 2131 because CISCO at least fails to disclose “a first multiserver platform” and “a second multiserver platform.” Further, CISCO fails to teach the identical invention in as complete detail as is contained in independent claim 9, as required by MPEP 2131 because CISCO at least fails to disclose “a first multiserver platform” and “a second multiserver platform.”

Further, the Applicant notes that “multiserver platform” is a known term in the art. As already stated above, CISCO does not disclose the use of a “multiserver platform.” Even a broad interpretation of CISCO cannot overcome at least this deficiency.

**H. Rejection of Dependent Claim 10**

Claim 10 depends on independent claim 9. Therefore, the Applicant submits that claim 10 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claim 9.

Accordingly, the Applicant submits that claim 10 is allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 10.

**I. Rejection of Dependent Claim 11**

Claim 11 depends on dependent claim 10, which depends on independent claim 9. Therefore, the Applicant submits that claim 11 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 9 and 10.

Accordingly, the Applicant submits that claim 11 is allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 11.

**J. Rejection of Dependent Claim 12**

Claim 12 depends on dependent claim 10, which depends on independent claim 9. Therefore, the Applicant submits that claim 12 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 9 and 10.

Accordingly, the Applicant submits that claim 12 is allowable over the reference cited in the Final Office Action at least for the above reasons. The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 12.

**K. Rejection of Dependent Claim 13**

Claim 13 depends on independent claim 9. Therefore, the Applicant submits that claim 13 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claim 9. The Applicant also submits that CISCO does not disclose or suggest at least the limitation of "further comprising at least one central switch coupled to at least said first switch blade of said first multiserver platform and said second switch blade of said second multiserver platform," as recited by the Applicant in claim 13.

With regard to claim 13, the Final Office Action states that CISCO discloses "further comprising at least one central switch (page 9, Figure 9, top central switch). See Final Office Action, Page 10, Lines 6-8. However, nowhere in CISCO is there any

mention of “a central switch.” Rather, the icon referred to by the Examiner in Figure 9 is a router, not a central switch as alleged by the Examiner. For example, CISCO’s Figure 2 shows the same icon and labels it “Cisco Router.” Additionally, CISCO’s Figure 7 shows the same icon and labels it “Cisco 7000.” The Applicant submits that a router is different than a central switch. Thus, CISCO fails to disclose “a central switch” as recited in Applicant’s dependent claim 13.

The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 13.

#### **L. Rejection of Dependent Claim 14**

Claim 14 depends on dependent claim 13, which depends on independent claim 9. Therefore, the Applicant submits that claim 14 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 9 and 13.

The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 14.

#### **M. Rejection of Dependent Claim 15**

Claim 15 depends on dependent claim 14, which depends on dependent claim 13, which depends on independent claim 9. Therefore, the Applicant submits that claim

15 is allowable over the reference cited in the Final Office Action at least for the reasons stated above with regard to claims 9, 13 and 14.

The Applicant also reserves the right to argue additional reasons beyond those set forth above to support the allowability of claim 15.

## CONCLUSION

For at least the foregoing reasons, the Applicant submits that claims 1-15 are in condition for allowance. Reversal of the Examiner's rejection and issuance of a patent on the application are therefore requested.

No fee is believed due with respect to this Revised Appeal Brief. **The fee for the Appeal Brief has already been paid.** See June 16, 2008 Appeal Brief. The Commissioner is authorized, however, to charge any additional fees or credit any overpayment to the deposit account of McAndrews, Held & Malloy, Ltd., Account No. 13-0017.

Respectfully submitted,

Date: 26-SEP-2008

By: /Philip Henry Sheridan/  
Philip Henry Sheridan  
Reg. No. 59,918  
Attorney for Applicant

McANDREWS, HELD & MALLOY, LTD.  
500 West Madison Street, 34th Floor  
Chicago, Illinois 60661  
(T) 312 775 8000  
(F) 312 775 8100

(PHS)

**CLAIMS APPENDIX**  
**(37 C.F.R. § 41.37(c)(1)(viii))**

1. A method for communication information in a server platform, the method comprising:

receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform;  
determining at least a server blade associated with a second multiserver platform for receiving at least a portion of said received at least one packet; and  
routing said at least a portion of said at least one received packet to at least said server blade.

2. The method according to claim 1, wherein said receiving further comprises receiving said at least one packet by at least one of a second switch blade associated with a third multiserver platform and a central switch.

3. The method according to claim 2, further comprising if said at least one packet is received by said central switch, communicating said at least a portion of said at least one received packet to at least a third switch blade associated with said second multiserver platform via at least one communication link that couples said central switch directly to said at least said third switch blade.

4. The method according to claim 1, further comprising processing said routed at least a portion of said at least one received packet by said at least said server blade.

5. A machine-readable storage having stored thereon, a computer program having at least one code section for communicating information in a server platform, the at least one code section being executable by a machine for causing the machine to perform steps comprising:

receiving at least one packet from at least one of a first switch blade associated with a first multiserver platform;

determining at least a server blade associated with a second multiserver platform for receiving at least a portion of said received at least one packet; and

routing said at least a portion of said at least one received packet to at least said server blade.

6. The machine-readable storage according to claim 5, further comprising code for receiving said at least one packet by at least one of a second switch blade associated with a third multiserver platform and a central switch.

7. The machine-readable storage according to claim 6, further comprising code for communicating said at least a portion of said at least one received packet to at least a third switch blade associated with said second multiserver platform via at least

one communication link that couples said central switch directly to said at least said third switch blade, if said at least one packet is received by said central switch.

8. The machine-readable storage according to claim 5, further comprising code for processing said routed at least a portion of said at least one received packet by said at least said server blade.

9. A system for communicating information in a server platform, the system comprising:

    a first multiserver platform comprising at least one of a network interface and a first switch blade; and

    at least a second multiserver platform comprising a second switch blade coupled to said first switch blade of said first multiserver platform.

10. The system according to claim 9, further comprising at least a third multiserver platform comprising a third switch blade coupled to at least one of said second switch blade of said second multiserver platform and said first switch blade of said first multiserver platform.

11. The system according to claim 10, wherein said first multiserver platform, said second multiserver platform and said third multiserver are coupled in a daisy-chain configuration.

12. The system according to claim 10, wherein said first multiserver platform, and said third multiserver platform communicate via said second multiserver platform.

13. The system according to claim 9, further comprising at least one central switch coupled to at least said first switch blade of said first multiserver platform and said second switch blade of said second multiserver platform.

14. A system according to claim 13, further comprising at least a third switch blade of a third multiserver platform coupled to said at least one central switch.

15. The system according to claim 14, wherein said first multiserver platform, said second multiserver platform and said third multiserver platform communicate via said central switch.

## **EVIDENCE APPENDIX**

### **(37 C.F.R. § 41.37(c)(1)(ix))**

(1) CISCO Systems, Virtual LAN Communications, <http://web.archive.org/web/19990209172148/cio.cisco.com/warp/public/614/13.html>, July 14, 1995 ("CISCO"), entered into record by Examiner Hari P. Kunamneni in the March 27, 2007 Office Action.

**RELATED PROCEEDINGS APPENDIX**  
**(37 C.F.R. § 41.37(c)(1)(x))**

The Appellant is unaware of any related appeals or interferences.

# Virtual LAN Communications

---

- [Statement of Direction: Cisco VLAN Roadmap](#)
- [White Paper: Cisco IOS VLAN Services](#)
- [Technology Brief: VLAN Interoperability](#)

---

## Introduction

Today's cost-effective, high-performance **local-area network (LAN) switches** offer users superior microsegmentation, low-latency packet forwarding, and increased bandwidth across the corporate backbone. LAN switches also can segment networks into logically defined virtual workgroups. This logical segmentation, commonly referred to as virtual LAN (VLAN) communication, offers a fundamental change in how LANs are designed, administered, and managed. While logical segmentation provides substantial benefits in LAN administration, security, and management of network broadcast activity across the enterprise, there are many components of VLAN solutions that must be considered prior to large scale VLAN deployment.

These additional VLAN components include high-performance switches that logically segment connected end stations, transport protocols that carry VLAN traffic across shared LAN and **Asynchronous Transfer Mode (ATM)** backbones, layer 3 routing solutions that extend VLAN communications between **workgroups**, system compatibility and **interoperability** with previously installed LAN systems, and network management solutions that offer centralized control, configuration, and traffic management functions. Figure 1 summarizes these concepts. All of these components are critical for enterprise-wide VLAN solutions, because they provide the scalability necessary for migrating from an installed base of shared LAN technologies to the new, emerging architecture of per-user switched communications.

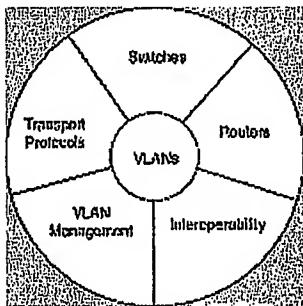


Figure 1: VLAN System Components

---

The first section of this document briefly discusses the importance of each one of these components within VLAN architectures. The second section reviews the benefits of VLANs and their applicability

within workgroups and across the enterprise backbone.

## Building VLAN Solutions

### Removing the Physical Boundaries

Conceptually, VLANs provide greater segmentation and organizational flexibility. VLAN technology allows network managers to group switch ports and users connected to them into logically defined communities of interest. These groupings can be coworkers within the same department, a cross-functional product team, or diverse users sharing the same network application or software (such as Lotus Notes users). Grouping these ports and users into communities of interest, referred to as VLAN organizations, can be accomplished within a single switch, or more powerfully, between connected switches within the enterprise. By grouping ports and users together across multiple switches, VLANs can span single building infrastructures, interconnected buildings, or even wide-area networks (WANs). VLANs completely remove the physical constraints of workgroup communications across the enterprise, as shown in Figure 2.

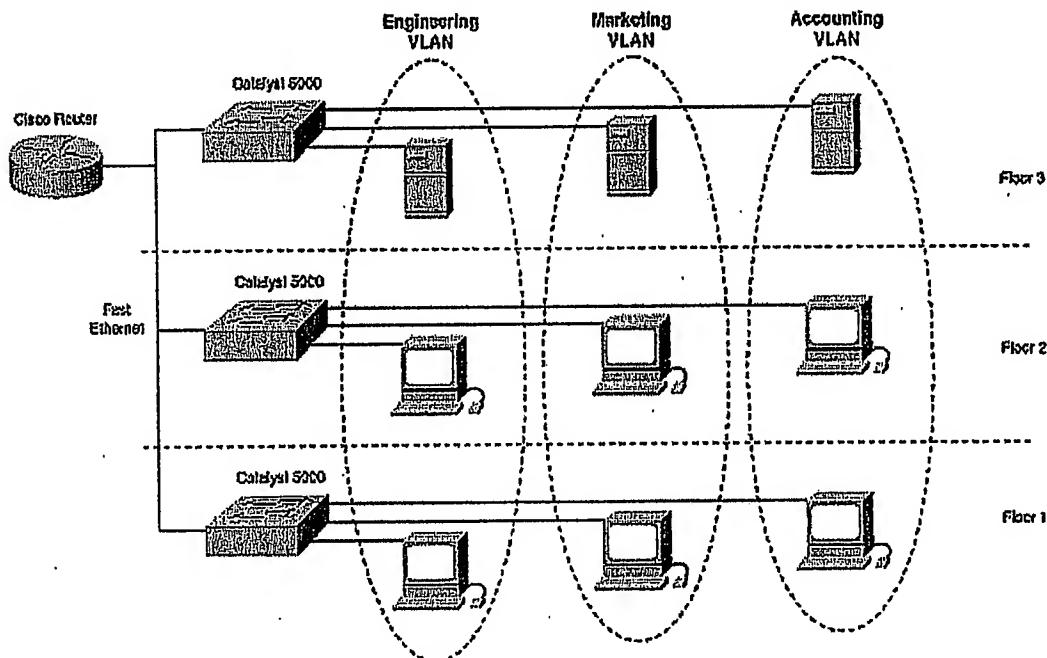


Figure 2: Logically Defined Networks (VLANs)

VLANs provide the ability for any organization to be physically dispersed throughout the company while maintaining its group identity. For example, accounting personnel can be located on the shop floor, in the research and development center, in the cash disbursement office, and in the corporate offices, while at the same time all members reside on the same virtual network, sharing traffic only with each other. Figure 3 illustrates a typical VLAN architecture that places these employees closer to their assigned areas of management and the people with whom they interact, while maintaining communication integrity within their respective organization.

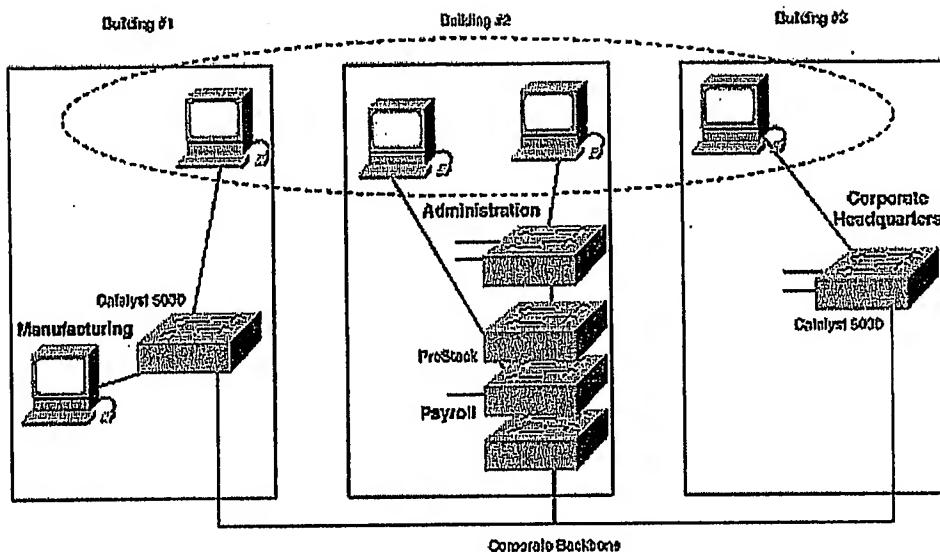


Figure 3: Logical Communication Between Users

Today's VLANs better match the way that companies are organized, and allow network managers to more closely align the network to the way that employees work and communicate.

### Switches -- the Core of VLANs

Switches are one of the core components of VLAN communications. They are the entry point for end-station devices into the switched fabric and for communication across the enterprise. Switches provide the intelligence to group users, ports, or logical addresses into common communities of interest. Each switch has the intelligence to make filtering and forwarding decisions by packet, based on VLAN metrics defined by network managers, and to communicate this information to other switches and routers within the network. And while today LAN switches are installed between shared segment hubs and routers located within the backbone, they will take on a larger, more significant role for VLAN segmentation and low-latency forwarding as they are deployed in the wiring closet. LAN switches offer significant increases in performance and dedicated bandwidth across the network, with the intelligence necessary for VLAN segmentation.

The most common approaches for logically grouping users into administratively defined VLANs include *packet filtering* and *packet identification*. Packet filtering is a technique that examines particular information about each packet based on user-defined offsets. Packet identification (tagging) uniquely assigns a user-defined ID to each packet. Both of these techniques examine the packet when it is either received or forwarded by the switch. Based on the set of rules defined by the administrator, these techniques determine where the packet is to be sent, filtered, and/or broadcast. These control mechanisms can be centrally administered (with network management software) yet are easily deployed throughout the network.

The concepts of packet filtering are very similar to those commonly used for routers. A filtering table is developed for each switch. This provides a high level of administrative control, because it can examine many attributes of each packet. Network managers can group users based upon MAC station addresses, network layer protocol types, and/or application types. Table entries are compared with the packets filtered by the switch. The switch takes the appropriate action based on the entries (see Figure 4).

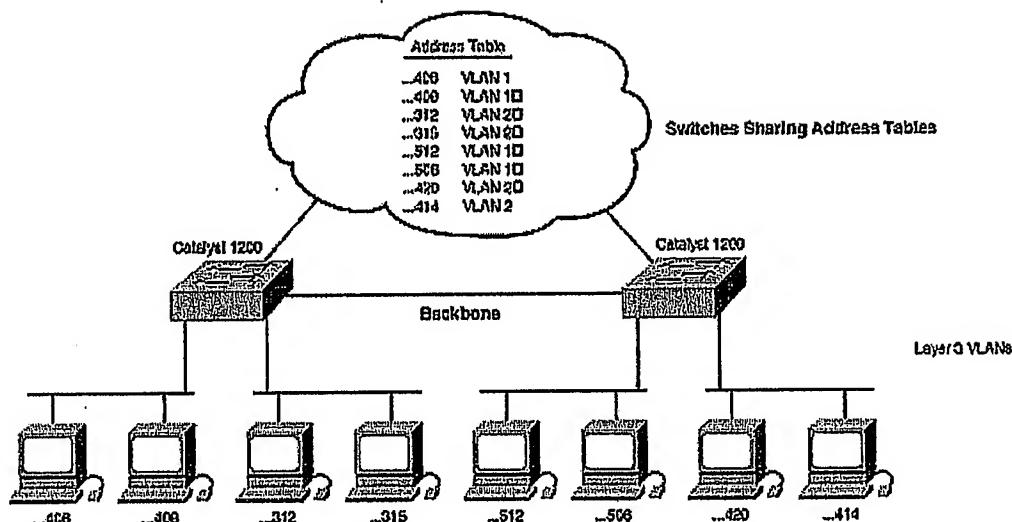


Figure 4: Packet Filtering

Packet filtering typically provides an additional layer of switch processing prior to forwarding each packet to another port or switch within the network, and it becomes more apparent as you filter deeper into each packet. This additional processing can effect switch latency and overall network performance. In addition, maintaining address tables adds an extra layer of administration per switch and requires synchronizing tables between switches.

VLAN packet identification (packet tagging) is a relatively new approach that has been specifically developed for switched communications. This approach places a unique identifier in the header of each packet as it is forwarded throughout the switch fabric. The identifier is understood and examined by each switch prior to any broadcasts or transmissions to other switches, routers, or end-station devices. When the packet exits the switch fabric, the switch removes the identifier before the packet is transmitted to the target end station. Over the past two years, packet identification has gained acceptance as switches have increased in popularity; packet identification functions at layer 2 and requires little processing or administrative overhead (see Figure 5).

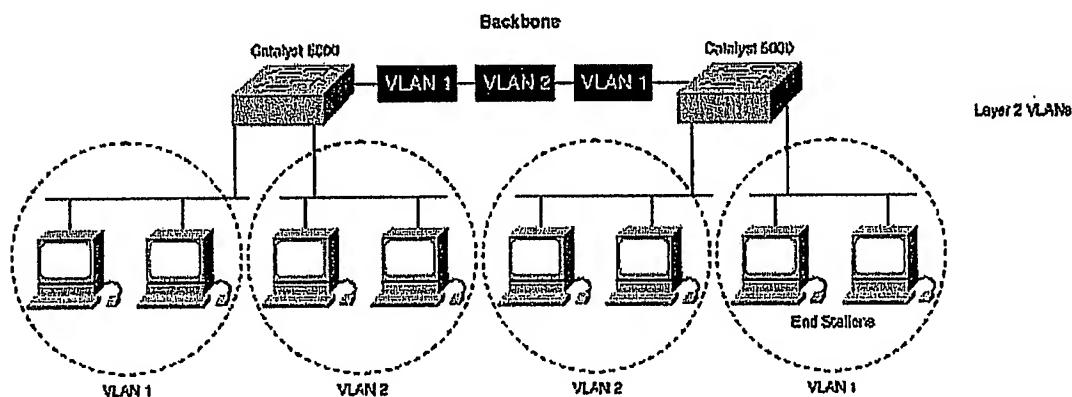


Figure 5: Packet Identification

The overall benefits of both approaches (packet filtering and packet identification) allow VLAN architectures that are nonintrusive to end-node applications and communication protocols. Switches provide all of the filtering, identification, and forwarding without any modification to the attached end station devices. This delivers a VLAN architecture that easily integrates with existing LAN applications while offering scalability and migration to ATM networks.

## Configuring VLANs

Users can be assigned to VLANs using several different configuration options that include static port assignments, dynamic port assignments, and multi-VLAN port assignments. These options are a function of the switch's capabilities (as mentioned in the previous section), the manner in which the stations are attached to each port on the switch, and the capabilities of the VLAN management software.

Stations directly attached to the ports on the switch provide the greatest flexibility for VLAN configuration and management. All stations can be uniquely assigned to VLANs. When they move to other physical locations using other directly attached switch port connections, they maintain their VLAN identities irrespective of their new locations. Stations connected to a switch through a shared hub are commonly grouped within the same VLAN because they all share the same switch port. While this approach is less flexible for each user on the network, it still provides highly desirable VLAN solutions for network managers. Additionally, hubs that provide multibackplane connection options increase the flexibility for unique VLAN assignments. Each backplane connection from the hub to a switch port can be individually assigned to a VLAN.

*Static* VLANs are ports on a switch that a network manager has statically assigned to a VLAN, using either a VLAN management application or has configured directly within the switch. These ports maintain their assigned VLAN configurations until the network manager takes another action. Although static VLANs require changes by the network operator, they are secure, easy to configure, and straightforward to monitor. These type of VLANs work well in networks where network moves are controlled and managed, where there is robust VLAN management software to configure the ports, and where network managers do not want to take on the additional overhead of maintaining end-station MAC addresses and custom filtering tables.

*Dynamic* VLANs are ports on a switch that can automatically determine their VLAN assignments with the aid of intelligent management software. Dynamic VLANs function based on their assignments to end-user station MAC addresses, logical addresses, or protocol type. These assignments are entered and maintained in a centralized VLAN management application. When a station is initially connected to an unassigned switch port, the appropriate switch checks the MAC address entry in the VLAN management database and dynamically configures the port with the corresponding VLAN configuration. The major benefits of this approach are less administration within the wiring closet when a user is added or moved, and centralized notification when an unrecognized user is added to the network. Typically, more administration is required up front to set up the database within the VLAN management software and to maintain an accurate database of all network users, as shown in Figure 6.

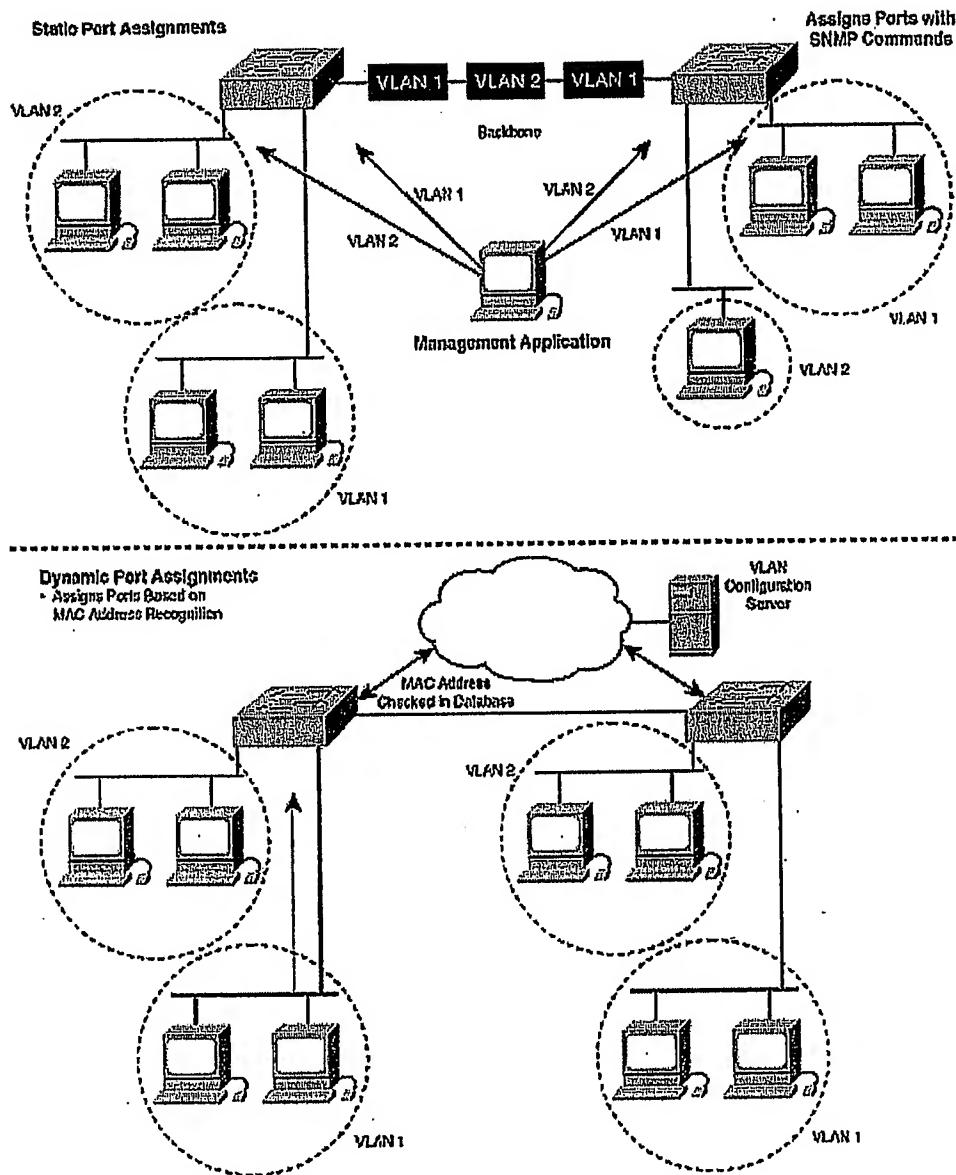


Figure 6: Configuring Ports to VLANs

Multi-VLAN port configurations provide communications among multiple VLANs concurrently either from a single port or a single user. This includes shared servers and users who need to belong to multiple workgroups. While there are several solutions on the market today that provide this functionality, there is an associated tradeoff. Concurrent port sharing across multiple groups dramatically reduces the firewalls between workgroups and the security these firewalls provide. These ports act as gateways into other VLAN groups and, in effect, create one larger VLAN. This approach does not scale well as the intersection between these VLAN groups becomes larger and larger.

For resources that need to participate in several VLANs concurrently, a better approach is to attach the end station directly to the backbone and to configure unique communication paths to each individual VLAN, thus providing resource sharing while maintaining the integrity of the VLAN firewalls. This approach has been defined in the ATM LAN Emulation draft standards and is also being evaluated for

implementation across shared-LAN backbones and switching architectures, illustrated by Figure 7.

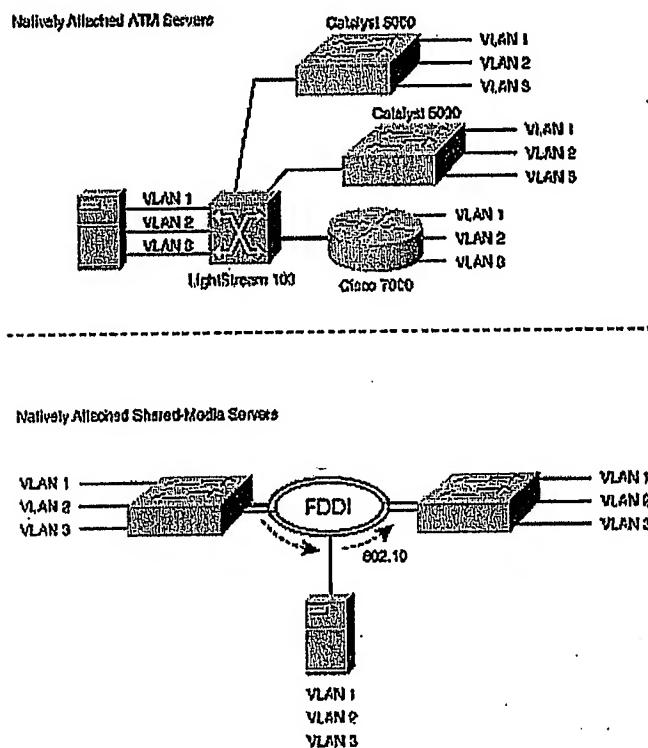


Figure 7: Servers as Part of Multiple VLANs

## Segmenting with Switching Architectures

Restructuring users according to logical associations across the enterprise rather than physical location is a fundamental shift away from the topologies deployed today. A large majority of networks currently installed provide very limited logical segmentation. Users are commonly grouped based on their connections into the shared hub and the router ports between these hubs. In addition, users on two different hubs segmented with a router cannot be connected to the same LAN segment. This topology provides segmentation only between the hubs, which are typically located on separate floors, and not between users connected to the same hub. It imposes physical constraints on the network and greatly limits the manner in which users can be grouped. And while some shared hub architectures provide a small degree of grouping capabilities, network managers are restricted in the way they can configure logically defined workgroups.

Switches remove the physical constraints imposed by a shared-hub architecture because they logically group users and ports across the enterprise. As a replacement for shared hubs, switches remove the physical barriers imposed within each wiring closet. Additionally, the role of the router evolves beyond the more traditional role of firewalls and broadcast suppression to policy-based control, broadcast management, and route processing and distribution. Equally as important, routers remain vital for switched architectures configured as VLANs because they provide the communication between logically defined workgroups (VLANs). Routers also provide VLAN access to shared resources such as servers and hosts, and connect to other parts of the network that are either logically segmented with the more traditional subnet approach or require access to remote sites across wide-area links. Layer 3

communication, either embedded in the switch or provided externally, is an integral part of any high-performance switching architecture.

External routers can be cost-effectively integrated into the switching architecture using one or multiple high-speed backbone connections. These are typically FDDI, Fast Ethernet, or ATM-type connections. These connections increase the throughput between switches and routers, provide a one-to-one logical association between the configured VLANs and layer 3 subnets, and consolidate the overall number of physical router ports required for communication between VLANs. As illustrated in Figure 8, this architecture not only provides logical segmentation, it greatly enhances the efficiency of the network.

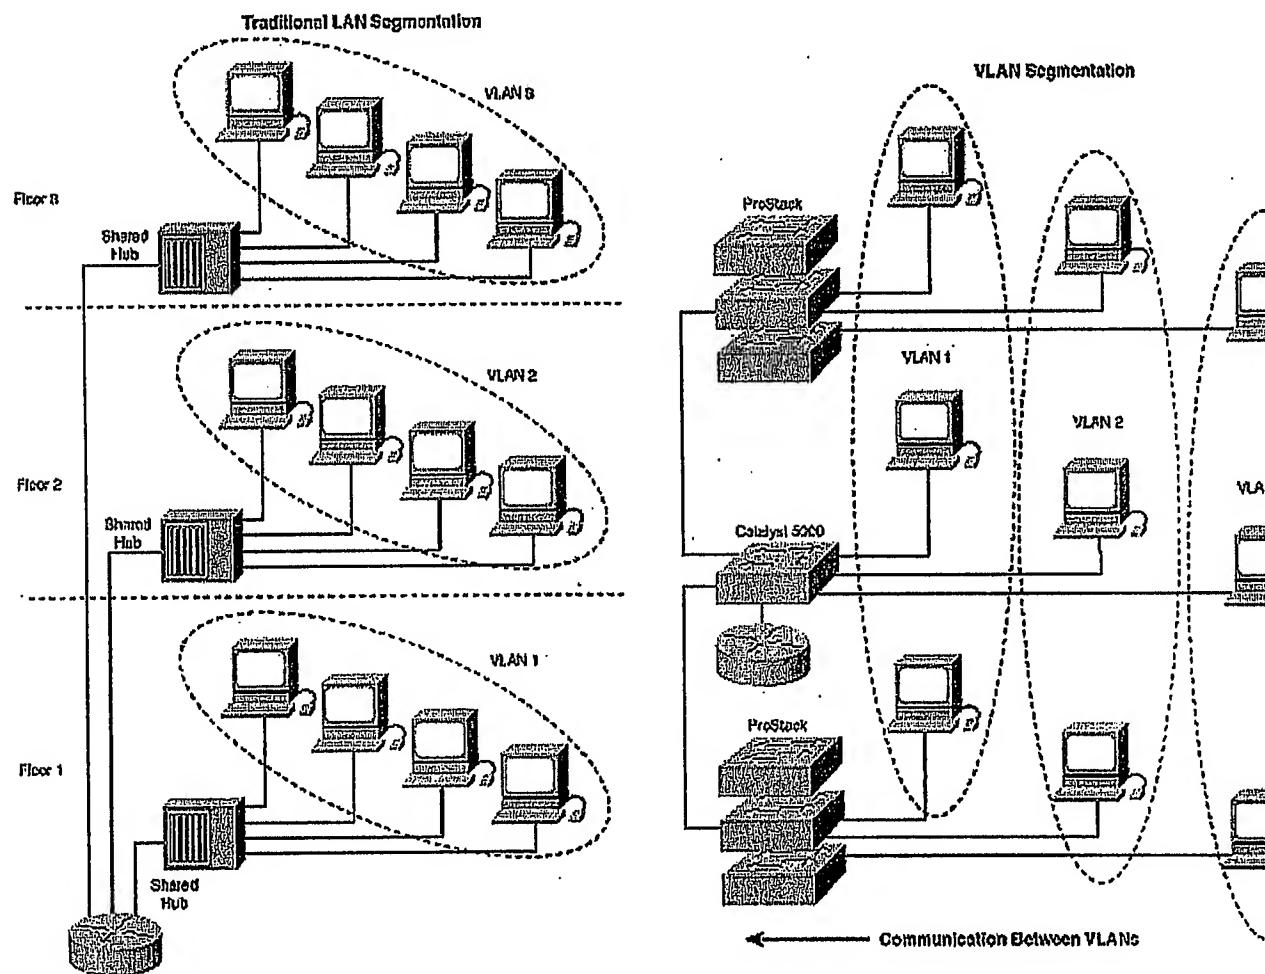


Figure 8: Topology Changes of VLANs

## VLANs across the Backbone

Important to any VLAN architecture is the ability to transport VLAN information between interconnected switches and routers that reside on the corporate backbone. It is the VLAN transport that enables enterprisewide VLAN communications. These transport capabilities remove the physical boundaries between users, increase the configuration flexibility of a VLAN solution when users move, and provide mechanisms for interoperability between backbone system components.

The backbone commonly acts as the aggregation point for large volumes of traffic. It also carries end-user VLAN information and identification between switches, routers, and directly attached servers. Within the backbone, high-bandwidth, high-capacity links are typically chosen to carry the traffic throughout the enterprise. The three most popular high bandwidth options include Fast Ethernet, Fiber Distributed Data Interface (FDDI), Copper Distributed Data Interface (CDDI), and ATM. Because switches and routers directly attach to the backbone, they must be able to transport VLAN information and interoperate with other network components.

In response to these requirements, several different transport mechanisms are being considered for communicating VLAN information across high-performance backbones. Among them is the LAN Emulation draft standard that has recently been approved by the ATM Forum and the IEEE 802.10 protocol which provides VLAN communication across shared backbones. Both of these define an interoperable mechanism for configuring and transporting VLANs across different backbone technologies.

The 802.10 proposal has been recommended by switching, routing, and hub vendors. Figure 9 shows typical applications for 802.10. This proposal defines a 32-bit addressing scheme within an 802.10 packet for VLAN identification, an addressing scheme nonintrusive to existing backbone architectures; however, it requires that switches include built-in software intelligence for enterprise VLAN communications. With the standardization of these two transport protocols, network managers can implement VLANs within individual workgroups, across the enterprise backbone, and gain access to WANs. In addition, Cisco has developed the inter-switch link (ISL) VLAN transport protocol to deliver efficient communication across Fast Ethernet backbones. Cisco will implement this as a de-facto standard and has made the specification available to vendors who want to interoperate.

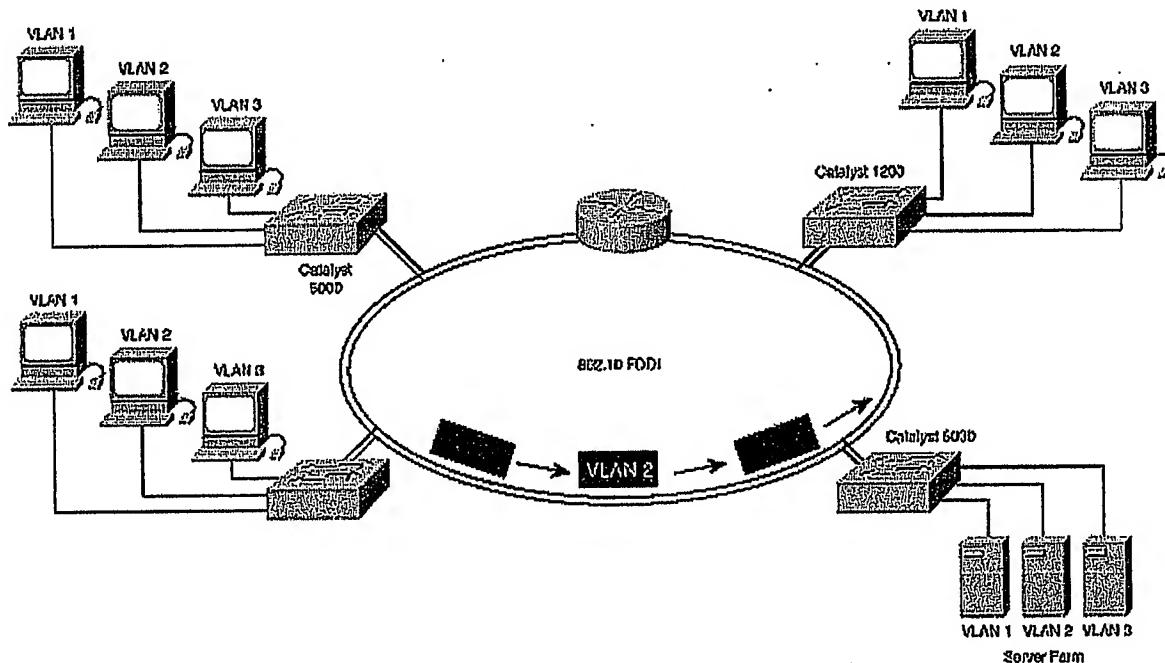


Figure 9: VLANs across FDDI Backbones

## VLAN Integration

Traditional network architectures are experiencing significant changes as they evolve toward greater microsegmentation, more capacity across the backbone, and dedicated circuit switching with the adoption of ATM. At the core of these changes are LAN-based switches with wiring closet applications, backbone switches for greater throughput performance, and ATM switches for dedicated circuit switching. As network managers migrate to these products, VLANs become a reality. Typically, the integration of VLANs begins with the first switch installation in a department or building. As the number of switches grows throughout the enterprise, VLANs become an enterprise-wide solution. These enterprise-wide VLANs require the transport mechanism, management tools, and layer 3 communication for logical segmentation and access across the network.

VLANs become a natural inclusion for LAN architectures as network designers and managers seek dedicated bandwidth to the desktop and segmentation based upon logical workgroups across the enterprise. Switching architectures that are VLAN-capable, along with routing solutions that interconnect VLANs, are evolutionary design changes compared with the physical segmentation that a majority of networks have in place today. VLANs are one of the essential technologies for breaking today's restrictive paradigm.

## **The Benefits of VLANs**

VLANs are often positioned as solving the problems associated with moves, adds, and changes. While they do reduce a large part of the administration costs when users change locations within a building or campus, VLAN technology provides many internetworking benefits that are equally as compelling. In addition to the reduced costs of administration, VLAN benefits include tighter network security with establishment of secure user groups, better management and control of broadcast activity, microsegmentation of the network without sacrificing scalability, load distribution of traffic across traffic-intensive switches ("hot spots" within the network), and the relocation of workgroup servers into secured, centralized locations.

### **Improved Administration Efficiencies**

Companies continuously reorganize as they seek productivity improvements. On average, between 20 and 40 percent of the workforce is physically moved every year. These moves, adds, and changes are one of a network manager's biggest headaches and one of the largest expenses relative to managing the network. Many moves require re-cabling, and almost all moves require new station addressing and hub and router reconfigurations. And, invariably, about the time managers stabilize their networks, more changes are requested.

VLANs provide an effective mechanism for controlling these changes and reducing much of the cost associated with hub and router reconfigurations. Users in a VLAN can share the same network "address space" regardless of their location. When users in a VLAN are moved from one location to another, as long as they remain within the same VLAN and are connected to a switch port, their network addresses do not change. Location changes can be as simple as plugging a user into a port on a VLAN-capable switch, or simply configuring the port on the switch to that VLAN, as shown in Figure 10. This greatly simplifies the rewiring, configuration, and debugging necessary to get the user back on line. It is a significant improvement over the techniques used within the wiring closet today. Moreover, router configuration remains intact; a simple move of a user from one location to another does not create any configuration modifications in the router as long as the user resides within the same VLAN.

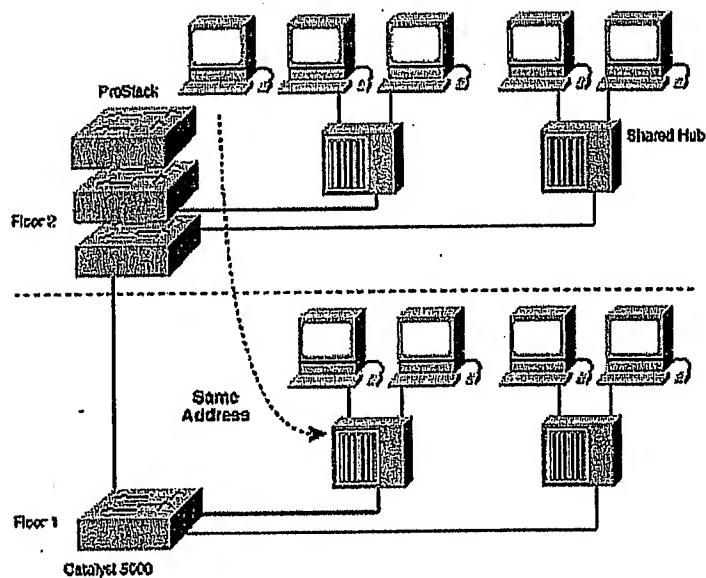


Figure 10: Controlling Broadcasts

### Controlling Broadcast Activity

Broadcast traffic, whether it is controlled through effective segmentation or by pruning an application's behavior, occurs in every network. Broadcast frequency depends on the types of applications, the types of servers, the amount of logical segmentation, and how these network resources are used. While applications have been fine-tuned over the last few years to reduce the number of broadcasts they send out, new multimedia applications are being developed that are broadcast- and multicast-intensive. Operationally, broadcasts can occur as a result of faulty network interface cards and communication devices. If not properly managed, they can seriously degrade network performance and can potentially bring down an entire network. This type of failure is primarily due to inadequate firewalls, internetworking loops, faulty network devices, or broadcast-intensive applications.

Network managers must take preventive measures to ensure against broadcast-related problems. One of the most effective measures is to properly segment the network with protective firewalls that minimize problems on one segment from damaging other parts of the network. Thus while one segment may exhibit excessive broadcast conditions as a result of a faulty network device or a mismanaged application, the rest of the network is protected with a firewall, commonly provided by a router. Firewall segmentation provides reliability, safeguards the network from the inefficient use of bandwidth, and minimizes the overhead of broadcast traffic allowing for greater throughput of application traffic.

As many designers migrate their networks toward switching architectures, they begin to lose the firewalls and safeguards that routers provide. By not placing any routers between the switches, broadcasts (layer 2 transmissions) are sent to every switched port. This is commonly referred to as a "flat" network where there is one broadcast domain across the entire network. The advantage of a flat switched network is that it provides very low latency and high throughput performance; the disadvantage is that it increases the susceptibility to broadcast traffic across all switches, ports, backbone links, and users.

Similar to routers, VLANs offer an effective mechanism for setting up firewalls within a switch fabric

and protecting the network against potentially dangerous broadcast problems. Additionally, VLANs maintain all of the performance benefits of switching. These firewalls are accomplished by assigning switch ports, and/or users to specific VLAN groups both within single switches and across multiple connected switches. Broadcast traffic within one VLAN is not transmitted outside the VLAN. Conversely, adjacent ports do not receive any of the broadcast traffic generated from other VLANs. This type of configuration substantially reduces the overall broadcast traffic, frees bandwidth for real user traffic, and lowers the overall vulnerability of the network to broadcast storms (see Figure 11).

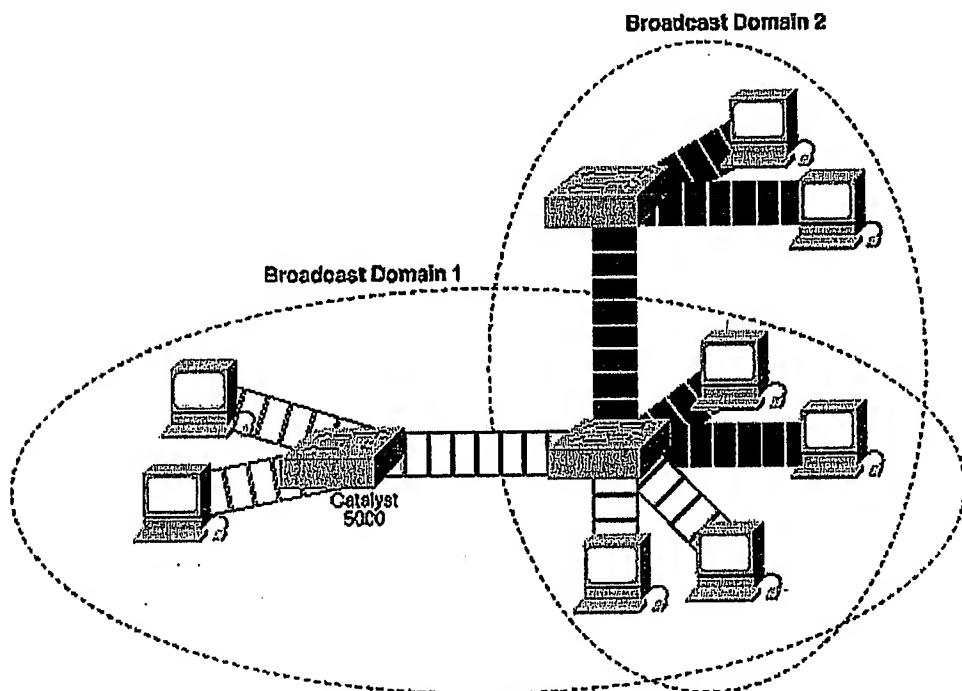


Figure 11: Added Security of Routers

Network managers can easily control the size of the broadcast domain by regulating the overall size of their VLANs, restricting the number of switch ports within a VLAN and the number of users residing on these ports. The smaller the VLAN group, the less effect broadcast traffic activity within the VLAN group has on everyone else within the network. Additionally, VLAN groups can be assigned based on the type of applications used and the amount of broadcasts these applications create. Users sharing an application that is very broadcast intensive are placed in the same VLAN group, while at the same time allowing the network manager to distribute the application across the campus.

### Enhanced Network Security

Over the past five years the use of LANs has increased exponentially. As a result, LANs often have confidential, mission-critical data moving across them. Confidential data requires security through access restriction. One of the inherent shortcomings of shared LANs is that they are relatively easy to penetrate. By plugging into a live port, an intrusive user has access to all broadcasts within the segment. The larger the broadcast group, the greater the access unless there are security control functions in the hub.

One of the most cost effective and easiest administrative techniques to increase security is to segment

the network into distinct broadcast groups. Additionally, it allows the network manager to restrict the number of users in a VLAN group and to disallow another user from joining without first receiving approval from the VLAN network management application. VLANs thus provide security firewalls, restrict individual user access, flag any unwanted intrusion to a network manager, and control the size and composition of the group.

Implementing this type of segmentation is relatively straightforward. Switch ports are grouped together based on the type of applications and access privileges. Restricted applications and resources are commonly placed in a secured VLAN group. Any users trying to tap into these secured VLANs are flagged by the network management software. Further security enhancements can be added using router access lists. These are especially useful when communicating between VLANs. On the secured VLAN, the router restricts access into the group as configured on both the switches and the routers. Restrictions can be placed based on station addresses, application types, protocols types, or even by time of day (see Figure 12).

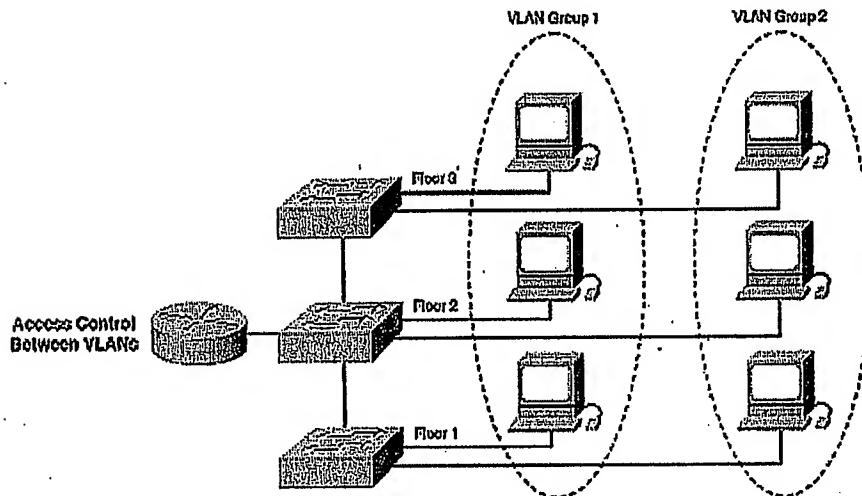


Figure 12: Simplification of Moves with VLANs

## Leveraging Legacy Hub Investments

Over the last three to five years, network administrators have installed a significant number of shared hub chassis, modules, and stackable devices. While many of these devices are being replaced with newer switching technologies as network applications require more dedicated bandwidth and performance directly to the desktop, these concentrators still perform useful functions in many existing installations. Network managers are leveraging their investments by connecting switches to the backplanes of the hubs. In the context of this discussion, a backplane hub connection defines any shared-media hub connection into a network backbone; stackable hubs, hub chassis, and even hub modules provide some form of this connection. It is the connections between the shared hubs and the switches that provide opportunities for VLAN segmentation. The greater the number of hub connections, the greater the opportunities for VLAN segmentation down to individual users.

Each hub segment (as defined within individual hub architectures) connected to a switch port can be assigned to a VLAN. Stations that share a hub segment are all assigned to the same VLAN group. If an individual station needs to be reassigned to another VLAN, the station will be relocated to the appropriate corresponding hub module. The interconnected switch fabric handles the communication

between the switching ports and automatically determines the appropriate receiving segments. The more the shared hub can be broken into smaller groups, the greater the microsegmentation and the greater the VLAN flexibility for assigning individual users to VLAN groups.

This furthers the migration to a high-performance switching architecture within enterprise LANs. With this approach, network managers can configure their shared hubs as part of the VLAN architecture and can share traffic and network resources that directly attach to switching ports with VLAN designations.

### **Centralized Administration Control**

Control of network broadcasts; planning moves, adds and changes; and establishing access privileges to the network and secured resources are common functions of the central planning and administration group. VLAN communications facilitate this type of planning by providing effective VLAN management applications that can be centrally configured, managed, and monitored.

## **Conclusion**

VLANs offer significant cost and performance benefits for a majority of the LANs installed today. These benefits are realized as network managers migrate to switched LAN architectures across the enterprise. And while VLANs are an integral part of ATM architectures, the concept and much of the technology has been designed into LAN-based switches that offer similar benefits across shared-LAN backbones. Further, end users' application need not change to realize these benefits. VLANs, as part of switching architecture, are invisible to end users. Finally, VLANs are more than simply a shared hub, routing, switching, or network management solution. It is the combination of all these components that provides powerful segmentation and efficient administration across the network.

---

Posted: Jul 14 15:32:01 1995

Copyright 1996 © Cisco Systems Inc.